Cyber Phenomenon Series

Cryptographic Identities for AI Agents

Scott Foote, Steve Foote

Last Updated: 27 March 2026

Phenomenati Consulting
www.phenomenati.com

6 Liberty Square, #2736
Boston, MA 02109
(508) 709-7990 (office)

CONFIDENTIALITY NOTICE: The contents of this document, including any attachments, are intended solely for stakeholders of Phenomenati Consulting, may contain confidential and/or privileged information, and are legally protected from disclosure.

Copyright © 2026 Phenomenati – All Rights Reserved.

<this page is intentionally blank>

Copyright © 2026 Phenomenati – All Rights Reserved.

Contents

Copyright © 2026 Phenomenati – All Rights Reserved.

1 Executive Summary

Over the last fifty years, identity management has expanded from local user accounts and passwords to a full discipline covering people, applications, devices, workloads, and federated digital ecosystems. The core questions, however, have remained remarkably stable: ***Who*** (or *What*) is acting? ***How*** do we know? ***What*** are they ***allowed*** to do? Under ***whose authority*** are they acting? And ***what evidence*** will survive an audit, a lawsuit, or an incident investigation? NIST's current digital identity guidance [NIST SP 800-63-4] continues to distinguish identity *proofing*, *authentication*, and *federation* as separate concerns, which is exactly the distinction enterprises now need for **AI Agents**.

AI Agents are not just another name for *service* accounts. They can operate semi-*autonomously*, invoke tools, traverse *trust* boundaries, act *on behalf* of specific principals, and produce *outcomes* that carry financial, regulatory, contractual, and reputational weight. NIST has now launched a dedicated AI Agent Standards Initiative, while the NCCoE has published a concept paper focused specifically on software and AI Agent identity and authorization, including questions around identification, authentication, proof of authority, least privilege, human oversight, tamper-proof logging, and non-repudiation [NIST AI Agent Standards Initiative; NIST NCCoE Concept Paper].

KEY CONCLUSIONS – WHAT DOES SUCCESS LOOK LIKE?

Jump forward to a time when you're running an investigation with legal actions arising...
What will you want to have in place for Agentic Identity, Authentication, Authorization, and an Audit trail?

- AI Agents should be treated as **first-class enterprise subjects** with their own registry, owners, lifecycle controls, and revocation paths - not as anonymous automation hidden behind shared secrets or generic service principals.
- Authentication is **necessary but insufficient**. Enterprises also need registration, sponsorship, software integrity evidence, runtime attestation, delegated authority, and policy-bounded sessions [NIST SP 800-63-4; RFC 9334].
- A practical cryptographic identity for an AI Agent should be **layered**:
  - a durable *registration* identity,
  - a short-lived attested *runtime* identity, and
  - an even shorter-lived *task* or *delegation* identity.
- High-impact actions should generate **signed, timestamped action receipts** plus provenance records that can be retained for audit, dispute resolution, and forensic analysis [RFC 3161; NIST SP 800-92].
- The strategic goal is not simply to authenticate an agent, but to make its actions **verifiable, attributable, policy-constrained, and defensible after the fact**.

> **Working Definition**
>
> A **cryptographic identity for an AI Agent** is a verifiable binding among
> (1) a stable agent identifier,
> (2) protected key material,
> (3) an issuing or sponsoring authority,
> (4) evidence about the software and runtime that currently holds the key,
> (5) explicit delegated authority and policy constraints, and
> (6) durable evidence of actions performed under that identity.

Copyright © 2026 Phenomenati – All Rights Reserved.

2 The Arc of Identity Management Across Five Decades

In practice, the history of identity management is the history of enterprises widening the set of ***actors*** they must ***trust*** and ***govern***. The subject began as the employee or system user at a terminal, then expanded to partners and customers, then to web applications, APIs, cloud workloads, and now autonomous software agents. Each wave has required better ways to bind identity claims to credentials, credentials to policy, and policy-based decisions to evidence.

Table 1: The Widening Identity Subject

Era	Primary Identity Subject	Control Breakthrough	Limitation Exposed by AI
1970s-1980s	Employees and local operators	Passwords, local ACLs, basic separation of duties	Identity stayed siloed and poorly portable across systems
1990s	Enterprise users and systems	Directories, network authentication, PKI, early role models	Trust often remained coarse-grained and perimeter-based [RFC 5280; NIST RBAC Model]
2000s	Employees, partners, customers	Federation, SSO, provisioning, compliance, electronic records and signatures	Identity assertions did not always travel cleanly into APIs and automation [15 U.S.C. 7001; EU's eIDAS]
2010s	Users, apps, APIs, mobile devices	OAuth / OIDC, MFA, WebAuthn, zero trust thinking	Service accounts and bearer tokens became too blunt for modern automation [OpenID Connect Core 1.0; WebAuthn Level 2; NIST SP 800-207]
2020s	Workloads, software agents, AI Agents	SPIFFE workload identity, attestation, verifiable credentials, AI risk governance	Now the enterprise must prove software identity, authority, and provenance at runtime [SPIFFE Overview; RFC 9334; W3C VC Data Model v2.0; NIST AI RMF 1.0]

If one step back defines the long arc, it is this: identity has become progressively less about a *login* event and more about **a chain of verifiable *trust***. By the time zero trust architecture formalized the principle that no implicit trust should be granted simply because an actor is on the network, identity had already become the primary input to every access decision [NIST SP 800-207].

AI Agents push that arc one step further. A modern agent may schedule meetings, initiate transactions, retrieve confidential data, synthesize recommendations, open tickets, write code, and invoke downstream agents. That is why NIST's 2026 standards work on agents explicitly highlights authentication and identity infrastructure as a dedicated research area [NIST AI Agent Standards Initiative].

Copyright © 2026 Phenomenati – All Rights Reserved.

3 Verifying Identity: Why Authentication Alone Is Not Enough

One of the most persistent sources of confusion in enterprise security is the tendency to collapse several distinct concepts into the single word – ***identity***. For practical governance… especially for AI Agents… the distinctions matter.

Table 2: What Each Discipline Answers

Discipline	Question it Answers
Identification / Registration	What ***subject*** is being named, sponsored, and recorded in enterprise governance?
Identity Proofing	What ***evidence*** supports the claim that this subject is who or what it purports to be? [NIST SP 800-63-4]
Authentication	Can the subject ***prove*** control of the credential or authenticator bound to that identity?
Authorization	What may this subject ***do*** in this context, against this resource, right now? [NIST SP 800-207]
Attribution	Can we ***prove later*** which subject performed a specific action?
Agency	On whose ***authority*** did the subject ***act***, within what delegated ***scope***, and with what ***approvals***?

For human users, proofing may rely on documentary evidence, biometrics, possession of devices, or in-person checks. For non-person entities such as services, workloads, or AI Agents, proofing is ***architectural*** rather than ***biometric***. The enterprise must establish sponsorship, ownership, software origin, deployment environment, key custody, operating purpose, and lifecycle controls. That is why simply authenticating with a secret or bearer token is a weak identity claim for an agentic system.

A service account can prove that it knows a secret; a signed JWT can prove that a token was minted by a trusted issuer; a workload certificate can prove control of a private key; a remote attestation result can prove that certain software or hardware claims were asserted about the runtime [RFC 7519; RFC 5280; RFC 9334]. None of those, by themselves, prove whether the actor was approved to execute a specific business decision or whether the resulting action should be *legally* or *operationally* ***attributable*** to the organization.

This is where AI changes the problem materially. An agent may begin with a legitimate task and then traverse tools, retrieved content, and downstream systems whose combined effect *exceeds* the original human *intent*. Prompt injection, retrieval poisoning, and tool confusion do not negate authentication, but they can sever the ***trust*** relationship between an authenticated agent and an authorized mission. The NCCoE concept paper correctly identifies proof of authority, conveyance of intent, human-in-the-loop controls, and tamper-proof logging as distinct design needs for software and AI Agents [NIST NCCoE Concept Paper].

COMMON FAILURE PATTERNS

- Using ***shared*** **API keys** for multiple automations, which destroys attributable identity.
- Treating a ***long-lived*** **service principal** as if it were sufficient for an AI Agent that changes tools, context, and runtime state.
- Allowing **delegated authority to *vanish* in transit**, so downstream systems see only a generic caller rather than the principal, purpose, and approval chain.
- Assuming that **logs alone** will reconstruct *intent*, even when the logs are mutable, incomplete, or detached from the cryptographic identity that executed the act.

Copyright © 2026 Phenomenati – All Rights Reserved.

4 How Verified Identities Are Represented Cryptographically

Cryptographic identity is best understood as a ***stack***, not a single *artifact*. A name without a key is not *trustworthy*; a key without governance is not *accountable*; and a signature without context is often not *enough* to support authorization, agency, or non-repudiation.

*Table 3: The Cryptographic Identity **Stack** for AI Agents*

Layer	Typical Artifacts	Why it Matters
Stable Identifier	Agent ID, service name, SPIFFE ID, DID	Provides persistent naming, governance, ownership, and lifecycle [SPIFFE Overview; W3C DID Core 1.0]
Protected Key Material	Private key in HSM, TPM, TEE, passkey authenticator	Provides proof that the claimant controls the key bound to the identity
Credential or Assertion	X.509 certificate, SVID, JWS/JWT, verifiable credential	Makes claims machine-verifiable by issuers and relying parties [RFC 5280; RFC 7515; RFC 7519; W3C VC Data Model v2.0]
Delegation Artifact	Short-lived token, token exchange result, approval assertion	Carries bounded authority, principal context, audience, and expiration [RFC 8693; OpenID Connect Core 1.0]
Evidence Artifact	Attestation evidence, signed receipt, trusted timestamp, append-only log entry	Supports audit, dispute resolution, non-repudiation, and forensics [RFC 3161; RFC 9162; RFC 9334]

Public key infrastructure remains the classical model for binding a subject name to a public key, with certificate path validation and revocation semantics defined in the X.509 PKI profile [RFC 5280]. That model still matters because certificates remain the backbone of TLS, device identity, workload identity, and many enterprise trust frameworks.

In web and API ecosystems, signed claims frequently travel through artifacts such as JSON Web Signatures and JSON Web Tokens. OpenID Connect layers identity semantics on top of OAuth 2.0 by standardizing signed ID Tokens that express claims about the authenticated subject [RFC 7515; RFC 7519; OpenID Connect Core 1.0]. For delegated and on-behalf-of patterns, token exchange provides a standard way to obtain a new security token appropriate to the downstream context [RFC 8693].

For human sponsors, phishing-resistant public-key credentials now have a mature standards base in WebAuthn and the broader passkey / FIDO ecosystem [WebAuthn Level 2]. That matters because every well-governed AI Agent still needs accountable human and organizational sponsors upstream of it.

For cloud-native workloads, SPIFFE popularized a simple but powerful pattern: *short-lived* cryptographic identity documents - SVIDs - issued to workloads so that services can authenticate each other using mTLS or JWT-based mechanisms [SPIFFE Overview]. For higher assurance, remote attestation architectures provide a way to *reason* about whether a runtime was in the expected state when a claim was made [RFC 9334].

Portable attribute frameworks such as verifiable credentials and decentralized identifiers provide another useful pattern: not every claim about an agent has to be embedded directly in a certificate. Some claims - such as *organizational role*, approved *tool classes*, or *regulatory status* - can travel as separately verifiable assertions under an enterprise trust framework [W3C VC Data Model v2.0; W3C DID Core 1.0].

Copyright © 2026 Phenomenati – All Rights Reserved.

A Critical Distinction

A **digital signature proves *control* of the private key at signing time**. It does ***not***, by itself, prove business *authority*, policy *compliance*, human *approval*, or legal *sufficiency*. Those additional meanings come from the surrounding **trust framework**, **delegated authority**, **approvals**, retained **evidence**, and the **controls** that govern key issuance and use.

Copyright © 2026 Phenomenati – All Rights Reserved.

5 A Practical Identity Model for AI Agents

Enterprises should resist the temptation to give an AI Agent one monolithic identity that is expected to cover design time, run time, delegation, and evidence. A more defensible model *separates* those concerns.

5.1 Registration Identity

The ***registration*** identity is the durable enterprise record for the agent. It should name the agent, its owner, its sponsoring organization, its business purpose, its approved tool categories, its data handling tier, and its lifecycle state. This is the record an auditor, architect, or incident responder should be able to inspect even when no runtime instance currently exists.

At minimum, the registration should capture a stable identifier, issuer, business owner, technical owner, approved model and agent versions, intended operating environments, escalation contacts, expiration or review date, and revocation status. NIST's emerging work on agent identity underscores the need to answer fundamental questions about what an agent is, how it is identified, and what metadata is required to manage it safely [NIST NCCoE Concept Paper; NIST AI Agent Standards Initiative].

5.2 Runtime Identity

The ***runtime*** identity is *ephemeral*. It belongs to a specific executing *instance* of the agent in a specific environment. It should be short-lived, minted only after the workload has authenticated and, where appropriate, been attested. This is where workload certificates, SVIDs, platform keys, and attestation evidence become powerful [SPIFFE Overview; RFC 9334].

The critical question is not only '*Which agent is this?*' but also '*Which exact instance, in which runtime, under which software state, is asking to act right now?*' If the answer cannot be bounded in time and environment, the enterprise is over-trusting the software.

5.3 Task or Delegation Identity

The ***task*** identity is shorter-lived still. It expresses what this instance is allowed to do for a particular principal, within a particular purpose, for a particular audience, and until a particular expiration. In effect, it is the cryptographic carrier of *agency*.

For example, an executive assistant agent might hold a durable registration identity and obtain a short-lived runtime certificate whenever it starts, but it should still require a purpose-bound task token to book travel, schedule meetings, approve routine expenses, or issue customer communications. The downstream relying party should be able to see not only the agent identity, but also the delegated principal, the authorized purpose, the risk level, and the time window.

> **Design Principle**
>
> For AI Agents, **one identity is rarely enough**. A durable registration identity answers ***who* is enrolled**. A short-lived runtime identity answers ***what* is executing now**. A purpose-bound delegation identity answers **what this execution is *allowed* to do, for *whom*, and for *how long***.

Copyright © 2026 Phenomenati – All Rights Reserved.

6 From Identity to Authorization, Access Control, and Agency

6.1 Authorization and Access Control

Role-based access control remains a durable foundation because enterprises still need clean mappings between business responsibilities and permissions [NIST RBAC Model]. But for AI Agents, RBAC alone is usually too coarse. A well-designed authorization decision should evaluate at least the subject identity, delegated principal, requested action, target resource, data sensitivity, runtime posture, model or agent version, and current approval state.

Zero Trust architecture makes the right demand here: ***authentication*** (AuthN) and ***authorization*** (AuthZ) are *discrete functions* performed *before* a session to an enterprise resource is established, and no implicit trust should be granted solely because the subject is on a particular network or within a legacy trust zone [NIST SP 800-207]. For agents, the most important policy boundary is often not the *network* edge but the ***tool boundary***.

In practical terms, every materially sensitive tool call - database query, message send, ticket closure, code merge, financial action, or document release - should be a policy enforcement point. The agent should present a short-lived, audience-restricted, purpose-restricted credential rather than a generic long-lived secret. Where delegated chains are required, token exchange and on-behalf-of patterns are far safer than reusing root credentials downstream [RFC 8693].

6.2 Legally Defensible Agency

The most important legal and governance question is usually *not* whether software is a "*legal person*". It is whether the enterprise can *prove* that the software acted as an *authorized instrument* of a natural or juridical person, within a defined scope, using reliable controls and tamper-evident records.

Electronic signature and trust service frameworks matter here because they establish that electronic signatures, contracts, and records generally should *not* be denied legal effect solely because they are electronic [15 U.S.C. 7001; EU's eIDAS]. But those laws do not remove the need for attribution, authority, consent, and trustworthy process. In other words, the signature *technology* is necessary, but the ***governance*** model is what makes the act defensible.

A high-assurance chain of agency for an AI Agent usually includes:

- a registered agent identity;
- a named organizational sponsor;
- phishing-resistant authentication by the human or system authorizing the task;
- a bounded delegation record;
- runtime assurance that the expected agent instance performed the action; and
- a signed, timestamped action receipt retained with the approval and policy decision [WebAuthn Level 2; RFC 3161].

For **high-impact acts**… such as payments, regulated notices, contract formation, code promotion, or irreversible data changes… many organizations will still want a *maker-checker* pattern, *dual* authorization, or a *human co-sign* requirement. Cryptographic identity should make such controls easier to enforce, not easier to bypass.

Copyright © 2026 Phenomenati – All Rights Reserved.

7 Decision Making, Provenance, Auditability, Non-Repudiation, and Forensics

7.1 Decision Provenance

Explainability and *Provenance* are related, but they are not the same thing. Explainability asks *why* a model or agent produced an output. Provenance asks *what entities*, activities, people, and systems *participated* in producing that output, and in what *relationship* [W3C PROV-DM]. For governance, Provenance is often the more operationally useful concept because it can be recorded, hashed, signed, and queried later.

NIST's Generative AI Profile explicitly highlights content provenance as a core topic alongside governance, pre-deployment testing, and incident disclosure [NIST AI 600-1]. For enterprise agents, a **provenance record** should be treated as **a first-class evidence artifact**, not as debug exhaust.

A MINIMUM PROVENANCE BUNDLE FOR CONSEQUENTIAL ACTIONS

- **Requesting principal** or source of delegation
- **Agent registration ID** and **runtime instance ID**
- **Model, prompt template, policy, and tool versions** used at the time of action
- **Hashes or immutable references** for prompts, retrieved content, and generated outputs where raw retention is impractical
- **Sequence of tool invocations**, policy decisions, approvals, and external responses
- **Digital signature and trusted timestamp** on the final action receipt [RFC 3161]

7.2 Auditing and Non-Repudiation

Traditional logging guidance remains highly relevant. NIST's log management guidance emphasizes that effective logging is a foundational security control, but the agent era requires one more step: the most important records should be *cryptographically bound* to the *identities* and *decisions* they describe [NIST SP 800-92].

Non-repudiation is often spoken about as if one control can 'deliver' it. In reality, it is a layered evidentiary posture. The enterprise needs secure identity issuance, strong authentication, key protection, well-ordered timestamps, integrity-protected logs, reliable retention, and governance over *who may authorize what*. Digital signatures and trusted timestamps are crucial, but they are strongest when combined with append-only transparency mechanisms and retained approvals [RFC 3161; RFC 9162].

Certificate Transparency offers a useful architectural analogy. The core idea is not limited to public TLS certificates: append-only logging can also be used for agent registrations, delegation grants, high-risk action receipts, and key lifecycle events so that material actions become observable, reviewable, and difficult to alter without detection [RFC 9162].

7.3 Forensic Analysis and Incident Response

If an AI Agent sends the wrong customer communication, misroutes a payment, leaks restricted data, or makes an unsafe operational change, the investigation will hinge on questions that conventional IAM logs often cannot answer:

- Which exact agent version acted?
- Which runtime held the key?
- Which prompt or retrieved content influenced the choice?
- Which human approved it?
- Which policies were evaluated?

Copyright © 2026 Phenomenati – All Rights Reserved.

- What downstream tools were invoked, in what sequence, and with what responses?

NIST's guidance on integrating forensic techniques into incident response remains directly relevant: forensic readiness should be built into system design, not added only after a serious event [NIST SP 800-86]. For AI Agents, that means preserving the *identity graph*, software bill of execution, attestation results, key lifecycle events, delegation chain, evidence hashes, and time-synchronized logs from the start.

Just as important, the **incident response** path must be operational: the enterprise should be able to ***quarantine*** an agent, ***revoke*** its active credentials, ***suspend*** new delegations, ***preserve*** evidence, and continue investigation without losing **chain of custody**.

Copyright © 2026 Phenomenati – All Rights Reserved.

8 Reference Architecture and Operating Model

A practical operating model for cryptographic AI identities can be understood as *three* interlocking *planes*: **identity**, **policy**, and **evidence**.

Table 4: Three Planes of Control

Plane	Core functions	Typical controls and artifacts
Identity Plane	Naming, registration, issuance, attestation, rotation, revocation	Agent registry, PKI, workload certificates, SVIDs, WebAuthn sponsor credentials, attestation evidence
Policy Plane	Delegation, authorization, step-up approval, risk scoring, separation of duties	PDP and PEP controls, short-lived task tokens, approval records, token exchange, tool-specific access policies
Evidence Plane	Receipts, provenance, timestamps, archival, transparency, investigation support	Signed action receipts, PROV records, RFC 3161 timestamps, append-only log, incident archive, retention controls

AN OPERATING SEQUENCE FOR HIGH-ASSURANCE AGENT ACTIONS

1. ***Register*** the agent, owner, purpose, tool classes, risk tier, and review date in a **controlled inventory**.
2. ***Authenticate*** the *sponsoring* human or upstream system with a phishing-resistant credential where required [WebAuthn Level 2].
3. ***Issue*** a *short-lived runtime credential* only after the workload authenticates and, where necessary, presents acceptable attestation evidence [SPIFFE Overview; RFC 9334].
4. ***Mint*** a *purpose-bound delegation token* for the specific task, including principal, audience, approved action classes, expiry, and approval state [RFC 8693].
5. ***Enforce*** authorization at *every* sensitive tool and data *boundary* rather than only at initial session establishment [NIST SP 800-207].
6. **Emit** a *signed action receipt*, ***bind*** it to a trusted *timestamp*, and ***store*** associated provenance and policy *records* [RFC 3161].
7. ***Rotate*** or ***revoke*** credentials *aggressively*, and ensure quarantine and investigation workflows are operational.

Copyright © 2026 Phenomenati – All Rights Reserved.

9 What Boards and Executive Teams Should Do Now

Most organizations do not need to invent new mathematics to govern AI Agents. They need to **apply mature identity and trust principles** with more rigor, better runtime specificity, and stronger evidence retention.

- Create an **authoritative registry** for all AI Agents and materially autonomous automations.
- Eliminate **shared secrets and generic service accounts** wherever agentic systems perform consequential work.
- Require a **layered identity model**:
  - *registration* identity,
  - *runtime* identity, and
  - *task-specific delegation* identity.
- Make the **tool boundary** the primary policy enforcement point for sensitive actions.
- Adopt **short-lived credentials** and explicit delegation rather than standing authority.
- Require **signed and timestamped receipts** for high-impact actions, especially those affecting money, regulated communications, code promotion, customer records, or safety-critical operations.
- Retain **provenance and audit artifacts** long enough to support incident response, legal holds, and independent review.
- Align the **CISO, CIO, CAIO, legal, privacy, records, and Internal Audit functions** around a shared accountability model rather than treating agent identity as only an engineering problem.

The **AI RMF** and the **GenAI Profile** both reinforce the need for governance structures that link risk management to concrete controls and evidence [NIST AI RMF 1.0; NIST AI 600-1]. Cryptographic identity is one of the few mechanisms that simultaneously serves *security*, *governance*, operational *resilience*, and defensible *accountability*.

Copyright © 2026 Phenomenati – All Rights Reserved.

10 Conclusion

The arc of identity management has always widened. It moved from human users to organizations, then to federated partners, then to applications and cloud workloads. AI Agents are the next widening of that circle… not because they are mystical new beings, but because they can now take *actions* that *matter* in the real economy.

The correct enterprise response is ***not*** to relax identity discipline in the name of innovation. It is to sharpen identity discipline so that every consequential agent action can be tied to a registered subject, an attested runtime, an explicit delegation, an evaluated policy decision, and durable evidence.

In that sense, cryptographic identity for AI Agents is both a control plane and an evidence plane. It determines who may act, and it determines what can be proven later. That is the foundation on which *trustworthy* enterprise agency will be built.

Copyright © 2026 Phenomenati – All Rights Reserved.

References

Adams, C., P. Cain, D. Pinkas, and R. Zuccherato. RFC 3161: *Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP).* IETF Proposed Standard, August 2001.
URL: https://www.rfc-editor.org/info/rfc3161

Birkholz, H., et al. RFC 9334: *Remote ATtestation procedureS (RATS) Architecture.* IETF Informational, January 2023.
URL: https://www.rfc-editor.org/info/rfc9334

Cooper, D., et al. RFC 5280: *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile.* IETF Proposed Standard, May 2008.
URL: https://www.rfc-editor.org/info/rfc5280

European Union. *Regulation (EU) No 910/2014 on Electronic Identification and Trust Services for Electronic Transactions in the Internal Market (eIDAS),* consolidated version current as of 18 October 2024.
URL: https://eur-lex.europa.eu/eli/reg/2014/910/2024-10-18/eng

Jones, M., et al. RFC 7515: *JSON Web Signature (JWS).* IETF Proposed Standard, May 2015.
URL: https://www.rfc-editor.org/info/rfc7515

Jones, M., et al. RFC 7519: *JSON Web Token (JWT).* IETF Proposed Standard, May 2015.
URL: https://www.rfc-editor.org/info/rfc7519

Jones, M., A. Nadalin, B. Campbell, J. Bradley, and C. Mortimore. RFC 8693: *OAuth 2.0 Token Exchange.* IETF, 2020.
URL: https://www.rfc-editor.org/info/rfc8693

Laurie, B., E. Messeri, and R. Stradling. RFC 9162: *Certificate Transparency Version 2.0.* IETF Experimental, December 2021.
URL: https://www.rfc-editor.org/info/rfc9162

National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* NIST AI 100-1, 2023.
URL: https://doi.org/10.6028/NIST.AI.100-1

National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile.* NIST AI 600-1, July 2024.
URL: https://doi.org/10.6028/NIST.AI.600-1

National Institute of Standards and Technology. *Digital Identity Guidelines.* NIST Special Publication 800-63-4, 2025.
URL: https://doi.org/10.6028/NIST.SP.800-63-4

National Institute of Standards and Technology. *SP 800-86: Guide to Integrating Forensic Techniques into Incident Response.* 2006.
URL: https://doi.org/10.6028/NIST.SP.800-86

National Institute of Standards and Technology. *SP 800-92: Guide to Computer Security Log Management.* 2006.
URL: https://doi.org/10.6028/NIST.SP.800-92

National Institute of Standards and Technology. *SP 800-207: Zero Trust Architecture.* 2020.
URL: https://doi.org/10.6028/NIST.SP.800-207

National Institute of Standards and Technology. *AI Agent Standards Initiative.* NIST CAISI program page, accessed March 2026.
URL: https://www.nist.gov/caisi/ai-agent-standards-initiative

National Institute of Standards and Technology, National Cybersecurity Center of Excellence. Booth, H., B. Fisher, R. Galluzzo, and J. Roberts. *Accelerating the Adoption of Software and AI Agent Identity and Authorization.* Draft concept paper, February 2026.
URL: https://www.nccoe.nist.gov/projects/software-and-ai-agent-identity-and-authorization

OpenID Foundation. Sakimura, N., J. Bradley, M. Jones, B. de Medeiros, and C. Mortimore. *OpenID Connect Core 1.0,* Final, 25 February 2014.
URL: https://openid.net/specs/openid-connect-core-1_0-final.html

Copyright © 2026 Phenomenati – All Rights Reserved.

Sandhu, R., D. Ferraiolo, and R. Kuhn. *The NIST Model for Role-Based Access Control: Towards a Unified Standard.* Proceedings of the Fifth ACM Workshop on Role-Based Access Control, 2000.
URL: https://doi.org/10.1145/344287.344301

SPIFFE Project. *SPIFFE Overview.* Project documentation, accessed March 2026.
URL: https://spiffe.io/docs/latest/spiffe-about/overview/

United States. 15 U.S.C. § 7001. *Electronic Signatures in Global and National Commerce Act* (E-SIGN), General Rule of Validity.
URL: https://www.govinfo.gov/link/uscode/15/7001

W3C. *Decentralized Identifiers (DIDs) v1.0: Core Architecture, Data Model, and Representations.* W3C Recommendation, 19 July 2022.
URL: https://www.w3.org/TR/did-core/

W3C. *PROV-DM: The PROV Data Model.* W3C Recommendation, 30 April 2013.
URL: https://www.w3.org/TR/prov-dm/

W3C. *Verifiable Credentials Data Model v2.0.* W3C Recommendation, 15 May 2025.
URL: https://www.w3.org/TR/vc-data-model-2.0/

W3C. *Web Authentication: An API for Accessing Public Key Credentials - Level 2.* W3C Recommendation.
URL: https://www.w3.org/TR/webauthn-2/

Copyright © 2026 Phenomenati – All Rights Reserved.